

IN THE CLAIMS

1. (Currently Amended) A method for providing a file system snapshot, comprising:
generating a snapshot dataset for a source file in a file system, wherein the snapshot dataset is contains substantially no data and no metadata empty; and
copying to a shadow inode into a first inode within the snapshot dataset, in response to only modifying metadata of the source file, at least a portion of metadata within a second an inode corresponding to the source file, when only metadata of the source file is modified, wherein a disk address of a data block corresponding to the source file is not copied to the shadow inode; and
storing, into the first inode, disk address values equal to a ditto address to indicate that the disk address is an invalid disk address.
2. (Currently Amended) The method of claim 1, further comprising:
copying to the shadow first inode in the snapshot dataset, in response to only appending to the source file, at least a portion of metadata within the second the inode corresponding to the source file, when the data block corresponding to the source file is only appended, wherein the disk address of the data block corresponding to the source file is not copied to the shadow inode; and
storing, into the first inode, disk address values equal to a ditto address to indicate that the disk address is an invalid disk address.
3. (Currently Amended) The method of claim 2, further comprising:
copying to the shadow first inode in the snapshot dataset the second inode corresponding to the source file and copying to the snapshot dataset the data block corresponding to the source file, when the data block corresponding to the source file is overwritten or deleted, wherein the shadow first inode includes a disk address of the data block which was written in the snapshot dataset.

4. (Currently Amended) The method of claim 3, further comprising:
accessing a shadow the first inode of the snapshot dataset corresponding to the a source file;
determining whether the first shadow inode includes a valid disk address,
wherein if the first shadow inode includes a valid disk address, then reading a data block referenced by the disk address; and
wherein if the shadow inode contains the ditto address~~does not include a~~ disk address, then retrieving the second an inode of the source file and retrieving a data block referenced by a disk address in the second inode of the source file.
5. (Currently Amended) The method of claim 3, further comprising:
copying to the first shadow inode in the snapshot dataset the metadata within the second inode corresponding to the source file and copying to the snapshot dataset an indirect block corresponding to the source file and at least one data block corresponding to the source file, when at least one of the data blocks corresponding to the source file are overwritten or deleted, wherein the first shadow inode includes a disk address of the indirect block which was written in the snapshot dataset and wherein the indirect block includes a disk address of at least one data block which was written in the snapshot dataset.
6. (Currently Amended) The method of claim 5, further comprising:
Accessing~~accessing~~ the first a shadow inode corresponding to the a source file;
determining whether the first shadow inode includes a valid disk address,
wherein if the first shadow inode includes a valid disk address, then retrieving an indirect block referenced by the valid disk address and at least one data block defined by at least one disk address in the indirect block; and
wherein if the first shadow inode does not include a valid disk address, retrieving an the second inode of the source file, then retrieving an indirect block referenced by a disk address in the second inode of the source file and retrieving at least one data block referenced by at least one disk address in the indirect block.
7. (Currently Amended) A system for providing a file system snapshot, comprising:

means for generating a snapshot dataset for a source file in a file system, wherein the snapshot dataset is contains substantially no data and no metadata empty; and

means for copying ~~to a shadow into a first inode within the snapshot dataset, in response to only modifying metadata of the source file, at least a portion of metadata within a second an inode corresponding to the source file, when only metadata of the source file is modified, wherein a disk address of a data block corresponding to the source file is not copied to the shadow inode; and~~

means for storing, into the first inode, disk address values equal to a ditto address to indicate that the disk address is an invalid disk address.

8. (Currently Amended) The system of claim 7, further comprising:

means for copying to the shadow first inode in the snapshot dataset, in response to only appending to the source file, at least a portion of metadata within the second the inode corresponding to the source file, when the data block corresponding to the source file is only appended, wherein the disk address of the data block corresponding to the source file is not copied to the shadow inode; and

storing, into the second inode, disk address values equal to a ditto address to indicate that the disk address is an invalid disk address.

9. (Currently Amended) The system of claim 8, further comprising:

means for copying to the shadow first inode in the snapshot dataset the second inode corresponding to the source file and copying to the snapshot dataset the data block corresponding to the source file, when the data block corresponding to the source file is overwritten or deleted, wherein the shadow first inode includes a disk address of the data block which was written in the snapshot dataset.

10. (Currently Amended) The system of claim 9, further comprising:

means for accessing a shadow first inode of the snapshot dataset corresponding to a the source file;

means for determining whether the shadow first inode includes a valid disk address,

means for reading, in response to a determination that the first inode contains a valid address, a data block referenced by the valid disk address; and

means for retrieving, in response to a determination that the first inode does not contain a valid address, an inode of the source file and retrieving a data block referenced by a disk address in the second inode of the source file.

11. (Currently Amended) The system of claim 9, further comprising:

means for copying to the shadow first inode in the snapshot dataset the second inode corresponding to the source file and copying to the snapshot dataset an indirect block corresponding to the source file and at least one data block corresponding to the source file, when at least one of the data blocks corresponding to the source file are overwritten or deleted, wherein the shadow first inode includes a disk address of the indirect block which was written in the snapshot dataset and wherein the indirect block includes a disk address of at least one data block which was written in the snapshot dataset.

12. (Currently Amended) The system of claim 11, further comprising:

means for accessing a shadow first inode corresponding to the a source file;

means for determining whether the shadow first inode includes a valid disk address,

means for retrieving an indirect block referenced by the valid disk address and at least one data block defined by at least one disk address in the indirect block; and

means for retrieving an the second inode of the source file, retrieving an indirect block referenced by a disk address in the second inode of the source file and retrieving at least one data block referenced by at least one disk address in the indirect block.

13. (Currently Amended) A computer readable medium including computer instructions for providing a file system snapshot, the computer instructions comprising instructions for:

generating a snapshot dataset for a source file in a file system, wherein the snapshot dataset is contains substantially no data and no metadata empty; and

copying to a shadow inode into a first inode within the snapshot dataset, in

response to only modifying metadata of the source file, at least a portion of metadata within a second an inode corresponding to the source file, when only metadata of the source file is modified, wherein a disk address of a data block corresponding to the source file is not copied to the shadow inode; and

storing, into the second inode, disk address values equal to a ditto address to indicate that the disk address is an invalid disk address.

14. (Currently Amended) The computer readable medium of claim 13, the computer instructions further comprising instructions for:

copying to the shadow first inode in the snapshot dataset, in response to only appending to the source file, at least a portion of metadata within the second the inode corresponding to the source file, when the data block corresponding to the source file is only appended, wherein the disk address of the data block corresponding to the source file is not copied to the shadow inode; and

storing, into the second inode, disk address values equal to a ditto address to indicate that the disk address is an invalid disk address.

15. (Currently Amended) The computer readable medium of claim 14, the computer instructions further comprising instructions for:

copying to the shadow first inode in the snapshot dataset the second inode corresponding to the source file and copying to the snapshot dataset the data block corresponding to the source file, when the data block corresponding to the source file is overwritten or deleted, wherein the shadow first inode includes a disk address of the data block which was written in the snapshot dataset.

16. (Currently Amended) The computer readable medium of claim 15, the computer instructions further comprising instructions for:

accessing a shadow the first inode of the snapshot dataset corresponding to the a source file;

determining whether the first shadow inode includes a valid disk address,

wherein if the first shadow inode includes a valid disk address, then reading a data block referenced by the disk address; and

wherein if the shadow inode contains the ditto address~~does not include a disk address~~, then retrieving the second an inode of the source file and retrieving a data block referenced by a disk address in the second inode of the source file.

17. (Currently Amended) The computer readable medium of claim 15, the computer instructions further comprising instructions for:

copying to the first shadow inode in the snapshot dataset the metadata within the second inode corresponding to the source file and copying to the snapshot dataset an indirect block corresponding to the source file and at least one data block corresponding to the source file, when at least one of the data blocks corresponding to the source file are overwritten or deleted, wherein the first shadow inode includes a disk address of the indirect block which was written in the snapshot dataset and wherein the indirect block includes a disk address of at least one data block which was written in the snapshot dataset.

18. (Currently Amended) The computer readable medium of claim 17, the computer instructions further comprising instructions for:

accessing a shadow-first inode corresponding to a source file;
determining whether the shadow-first inode includes a valid disk address,
wherein if the shadow-first inode includes a valid disk address, then retrieving an indirect block referenced by the valid disk address and at least one data block defined by at least one valid disk address in the indirect block; and

wherein if the shadow-first inode does not include a valid disk address, then retrieving an a second inode of the source file, retrieving an indirect block referenced by a disk address in the second inode of the source file and retrieving at least one data block referenced by at least one disk address in the indirect block.

19. (Currently Amended) A system for providing a file system snapshot, comprising:

a snapshot dataset for a source file in a file system, wherein the snapshot dataset is substantially empty; and

a shadow first inode in the snapshot dataset, the shadow first inode comprising metadata copied from an second inode corresponding to the source file, wherein the

~~shadow first inode is generated only when metadata of the source file is modified and wherein a disk ditto address is stored of a data block corresponding to the source file is not included in the shadow first inode.~~

20. (Currently Amended) The system of claim 19, further comprising: wherein in the first inode

~~a shadow inode in the snapshot dataset, the shadow metadata from the first inode is copied from an the second inode corresponding to the source file, wherein the shadow-first inode is generated only when the data block corresponding to the source file is appended and wherein the disk ditto address of the data block corresponding to the source file is not included is inserted into the shadow-first inode.~~

21. (Currently Amended) The system of claim 20, further comprising:

a data block corresponding to the source file in the snapshot dataset, wherein the data block is copied to the snapshot dataset when the original data block is overwritten; and

~~a shadow first inode in the snapshot dataset, the shadow first inode containing metadata copied from an inode corresponding to in the source file, wherein the shadow first inode is generated when the data block corresponding to the source file is overwritten or deleted and wherein the shadow first inode includes a disk address of the data block which was written in the snapshot dataset.~~

22. (Currently Amended) The system of claim 21, further comprising:

~~a shadow first inode in a snapshot dataset, the first inode corresponding to a data block within a source file;~~

~~a disk ditto address value stored included in the shadow first inode to indicate an invalid disk address; and~~

~~a data block referenced by the disk address;~~

~~an inode of the source file referencing the data block; and~~

~~a data block referenced by a disk address in the inode of the source file.~~

23. (Currently Amended) The system of claim 21, further comprising:

a shadow first inode in a snapshot dataset, the first inode corresponding to an indirect block within a source file;

a disk ditto address value stored included in the shadow first inode to indicate an invalid disk address; and

an indirect block referenced by the disk address;
at least one data block defined by at least one disk address in the indirect block;
an inode of the source file referencing the indirect block;
an indirect block referenced by a disk address in the inode of the source file; and
at least one data block referenced by at least one disk address in the indirect block.

24. (Currently Amended) A method for deleting a first snapshot of a file system, comprising:

determining the existence of an older snapshot;

wherein if there is an older snapshot, determining the existence of a ditto address reference in the older snapshot to an inode or a data block in the first snapshot, wherein the ditto address indicates an invalid disk address; and

wherein if there is no older snapshot, deleting any inode or data block in the first snapshot.

25. (Currently Amended) The method of claim 24, further comprising:

wherein if there is a reference ditto address in the older snapshot to an inode or a data block in the first snapshot, copying to the older snapshot the metadata in an inode or data block of an inode in the first snapshot being referenced and deleting any inode or data block in the first snapshot; and

wherein if there is no reference ditto address in the older snapshot to an inode or a data block in the first snapshot, deleting any inode or data block in the first snapshot.

26. (Currently Amended) A method for restoring a first snapshot of a file system, comprising:

wherein if there is a most recent snapshot, the most recent snapshot not being the first snapshot, copying to the most recent snapshot any inode or data block in the

file system referenced by the most recent snapshot, which shall be modified by the restoration of the first snapshot;

wherein if there is an inode or a data block in the first snapshot, copying the inode or data block in the first snapshot to the file system; and

wherein if there is a ditto disk address in the first snapshot, wherein the ditto address indicates an invalid disk address, copying to the filesystem the inode or data block of the most recent snapshot that corresponds to an inode with referenced by the ditto disk address and that contains a valid disk address to the file system.